

proportion. Hence you would never prefer both I to II and IV to III. So if you do have these preferences, you aren't doing that. What else is there to do?

Decision Criteria for "Partial Ignorance"

There are, at least, the four criteria considered earlier in this chapter for "games against Nature"; we must reconsider them in this new context, interpreting the relevant events as proportions of Yellow (Black) balls. These criteria do not allow for any personal discrimination between the relative probabilities of events considered possible, and three of them (excluding the Laplace rule) conflict with the inference of any relative personal probabilities for these events. Therefore, if we admitted all proportions between 0 and 1 as "possible," all four of these criteria would be inappropriate for the same reason as before; virtually all subjects regard, and act "as if," proportions of Yellow (Black) balls greater than $2/3$ were "less probable": indeed, "impossible" ("null"). But if we restrict attention to proportions between 0 and $2/3$, on the basis of the information supplied in the experiment, the most obvious inadequacies of these rules vanish.

In his analysis of these criteria, Milnor suggests parenthetically (but does not follow up):

Our basic assumption that the player has absolutely no information about Nature may seem too restrictive. However, such no-information games may be used as a normal form for a wider class of games in which certain types of partial information are allowed. For example, if the information consists of bounds for the probabilities of the various states of Nature, then by considering only those mixed strategies for Nature which satisfy these bounds, we construct a new game having no information.¹

¹Milnor, op. cit., p. 49.

We shall find no use for the ill-defined notion, "absolutely no information," nor for the related, subjective concept, "complete ignorance," except as the latter is defined behaviorally by the two conditions on preferences specified earlier. Nor have we yet defined any operational basis for a meaningful, general concept of "information consisting of bounds for probabilities" (see Chapter Eight). But in this particular example the latter concept does have an obvious interpretation; on the assumption that "knowledge" of a given proportion of colors in the urn would lead to betting behavior consistent with a definite probability distribution over events, we can speak of information on the bounds of "possible" proportions as information conveying bounds on the probabilities of various states of Nature. This is, of course, a special, considerably artificial situation; as Milnor goes on to say: "Unfortunately in practice partial information often occurs in vague, non-mathematical forms which are difficult to handle." In Chapter Eight we shall consider the highly developed analyses of B. O. Koopman and I. J. Good which give quantitative expression to partial information in its more general forms; but we shall then find that the present, simple example provides a more suitable model for the general case than it might appear.

It will simplify discussion if we modify the example to allow only a small number of distinct, alternative "possible" proportions of Yellow balls in the urn: each proportion determining a distinct, "definite" probability distribution over the events Red, Yellow, Black in a single drawing from the urn. Let us imagine that the urn is known to contain not 90 balls but only three, of which one and only one is Red. You are informed that three possibilities for the composition of the remaining two balls are exhaustive:

anxious to be released from a commitment to 1 or 2 after the "super-lottery" had been conducted. This attitude is closely related to another aspect of my preferences to which Pratt does not refer but which he would, no doubt, find equally paradoxical. In saying that I would pay the same amount for one compound lottery as for the other, I am saying, in particular, that I ^{might} ~~would~~ pay more for the lottery ticket $(1, 2; \frac{1}{2})$ than I ^{might} ~~would~~ pay for either of its "prizes," the pure strategies 1 or 2, if either of these had become available to me through circumstances other than the operation of a "stochastic process" of precisely these ^{assumed} characteristics (i.e., involving this set of "prizes" and these "known" probabilities). And if action 1 (or action 2) had become available to me as a result of this particular process, I ^{might} ~~would~~ pay more for it than if it had become available through some other, undefined process.

This pattern of preferences has nothing essentially to do with "suspicion" of the unknown circumstances that ^{might} have led to the offer of action 1. It could hold even though I were convinced that the person offering me these bets had no pecuniary interest in the outcome of a given bet; that he had, for example, only an unprejudiced interest in determining my actual expectations and degrees of uncertainty, and that he had no more and perhaps much less information than I on the actual outcomes to be expected from the bets. The point is not that I distrust the nature of the process that ^{might} have led to an offer of action 1 in isolation, but that I simply do not know what it was; whereas if 1 became available to me as the result of a compound lottery $(1, 2; \frac{1}{2})$ I know that the same process provided an "equal probability" of giving me the opportunity to choose action 2. It does not follow necessarily, nor in every case, that this

4.) If, on the contrary, I committed myself to 1 and 2 before flipping, Pratt argues, then if the coin came up heads I would be willing to pay a dollar (say) to be released from my commitment to 1 so that I could choose the action I really prefer, 3; and likewise, if the coin came up tails I would be willing to pay a dollar to be able to choose 4 instead of 2. Why not save the dollar by picking 3 and 4 in the first place: even if I had to pay a small premium for this option?

The flaw in this argument is that, on the basis of my own introspection and analysis of the problem, I would not be willing to pay a dollar, or anything, to be released from a commitment to 1 and 2 after flipping. The reasoning that leads me to this conclusion is not self-evident, which is just why the somewhat duplicative discussion of this example seems worthwhile. It remains true that if confronted with a choice between actions 1 and 3, I would ^{considerable} prefer 3 to 1 and might pay a premium of a dollar to get it: if these were the only alternatives available and I ^{did not know enough about the process and knew too much that was contradictory about the process that} were ignorant of the stochastic properties of the process which had made this particular pair of alternatives available instead of some ^{to assign definite probabilities to all possible offers.} Yet I maintain that I would be indifferent between committing myself to 1 or to 3 in advance of the particular compound lottery that Pratt describes, and indifferent between these two actions after flipping, in the event that the ^{given} lottery presented me with this particular pair of actions and I were permitted to choose freely after all.

delete italics.

My rationale for indifference between the two compound lottery tickets $(1, 2; \frac{1}{2})$ and $(3, 4; \frac{1}{2})$ is perhaps sufficiently clear; it is essentially the same one that Pratt gives. What may be less clear is why I would not be

does not seem reasonable to me to make a decision you will soon regret (after flipping), so you seem to me to be forced to choose 3 and 4. Now this choice gives you an 'unambiguous' probability $1/4$ of obtaining the preferred prize a, and the opposite choice (1 and 2) does also. It certainly seems unreasonable to me to prefer (3,4) to (1,2) when they give identical, 'unambiguous' probability distributions over the consequences. (My italics.)

The example here is, of course, precisely the one that we have been discussing above. Again we consider the choice between the "compound lottery tickets" $(1,2;\frac{1}{2})$ and $(3,4;\frac{1}{2})$ where 3 and 4 are actions with

~~un~~ambiguous distributions of payoff, and where, considered as pure alternatives, 3 is preferred to 1 and 4 is preferred to 2. It will be clear from our earlier discussion that I agree thoroughly with Pratt's assertion that the two "mixed strategies" give "identical, 'unambiguous' probability distributions over the consequences," and hence that, like Pratt, I see no reasonable basis for preferring one to the other. But I do not agree that indifference between these two mixed strategies is necessarily inconsistent with the earlier pattern of preferences between the alternative pure actions. Let us examine closely Pratt's reasons for thinking otherwise.

Pratt's key assumption is in the first sentence italicized above: that after flipping a coin to determine the pair of actions available to me, I would, if the coin came up heads, ^{prefer to} choose 3 over 1, and if the coin came up tails, ^{prefer to} choose 4 over 2. He concludes from this that if I had to commit myself to a choice from each pair before knowing which pair would become available as a result of the coin toss, I would be "forced to choose 3 and 4," because to choose instead 1 and 2 would be to make a decision I would surely "soon regret (after flipping)." (By Pratt's logic, a choice of 1 and 4 or of 2 and 3 would expose me to a 50% chance of "regret after flipping" with no compensating chance of gain over the choice of 3 and

longer be available. If they tend to act "as if" the outcomes of "a bet on Red_I" were ambiguous, no amount of coin-flipping should serve to lessen this tendency, for it cannot remove the ambiguity.

Anyone who uses this rule, then, will violate the von Neumann-Morgenstern utility axioms in his choices between certain prospects that offer ambiguous bets as prizes. Is this a powerful argument for rejecting the reasonableness of his behavior?

At least two proponents of the Savage approach have argued that it is. Although we have already covered the substance of their comments above, it is worthwhile at the cost of some duplication to examine their arguments in detail, since ~~they appear to throw~~ ^(their language appears) the actions of our hypothetical violator into a more questionable light than might appear from my discussion.

John Raiffa/Pratt Criticism

John Pratt has focused his comments upon the following example from Chapter IV:¹

	100		50	50
	R _I	B _I	R _{II}	B _{II}
1.	a	b	b	b
2.	b	a	b	b
3	b	b	a	b
4	b	b	b	a

If there are 50 each of R_{II} and B_{II} and 100 I's divided in an unknown way, you would prefer 3 to 1 and 4 to 2, and you claim this is not unreasonable. Suppose you are now offered the following opportunity: you may flip an 'unambiguously' fair coin (or better, do the equivalent with another urn); if you get heads, you may choose 1 or 3; if tails, you may choose 2 or 4. If you make your choice after flipping, you will surely choose 3 or 4.
If you are required to choose before flipping, it certainly

¹Private communication; in this case, "you" is me.

Raiffa has conducted such hypothetical experiments with many classes of students at the Harvard Business School and graduate classes in statistics. He reports that a majority of students will offer considerably more ^{initially} _A for the opportunity to choose between actions I and II than for the opportunity to choose between actions III and IV: e.g., \$35 in the first instance and \$5 in the second. However, when the availability of a "mixed strategy" based on coin-flipping is brought to their attention, he asserts that eventually they change their opinions and offer as much for the second opportunity as for the first. As we have seen, this conclusion does not contradict the hypothesis that they follow our decision rule concerning ambiguity. If they do obey the Savage axioms with respect to some hypothetical coins and urns with known, 50-50 composition, then they should indeed conclude, after analysis, that these various options are equivalent.

By the same token, it is important to notice that such experiments are crucially different from those suggested in Chapter IV, in which you were never invited to choose both a color and an urn freely; in those, randomized strategies with the effect of banishing ambiguity were infeasible. It would be interesting to know how much students who had finally been persuaded by Raiffa to offer as much for a free choice between III and IV as for a free choice between I and II would subsequently be willing to pay for, say, action III alone, offered in isolation as the sole option. (Raiffa conducted no experiments in which the students were not offered free choice of color.) My conjecture is that they might well go back to their original, low bids, e.g., \$5. It is clear that the rationale used effectively to influence their choices in the earlier situation would no

are ambiguous and the ^{latter} former, subjectively, is not. It is entirely consistent with a pattern of "penalizing" ambiguity ^{of payoff probabilities} to prefer the lottery ticket, or "mixed strategy," to either of these two ambiguous components.

It is something of a coincidence that the payoff distribution to this particular prospect is unambiguous where the component strategies are not, and this fact may not be immediately obvious to many subjects.

However, the following argument, ^{suggested by Howard Raiffa,} ~~addressed to them~~, may clarify the situation. ^{for them}

Imagine that the ball has already been drawn from Urn I, but you have not yet seen its color. If, on the one hand, it is actually red, then the prospect (III, IV; $\frac{1}{2}$) offers you an unambiguous 50-50 chance of \$100 or \$0; for if the coin turns up Heads, you choose action III and collect \$100 when the ball is revealed, and if it is Tails, you choose action IV and collect nothing. Likewise if the ball is black; the

prospect offers you exactly a 50-50 chance of III (paying \$0 in this case) or IV (paying \$100). So you are "guaranteed" ^{exactly} a 50-50 chance of \$100 or

\$0 "whether the ball is red or black" (^{assuming the probabilities of Heads or Tails to be unambiguous, or 'definite', for you).}

~~Howard~~ Raiffa ~~has~~ pointed out that you should be willing to pay at least as much as you would pay for this prospect, or for actions I or II, with respect to Urn II, for the opportunity to choose freely between actions III and IV with respect to Urn I. For given this opportunity, nothing can prevent you from flipping a coin between the two actions, choosing III if heads, IV if tails. In other words, if you are allowed to choose either III or IV, you are also free to choose (III, IV; $\frac{1}{2}$), so the opportunity should be worth at least as much as the latter option. Σ

¹ Howard Raiffa, "Risk, Ambiguity, and the Savage Axioms: Comment," Quarterly Journal of Economics, Vol. LXXV, No. 4, November 1961, p. 693.

Insert A

The probability of any one of these joint events will depend on the probability p assigned to a drawing of Red from Urn I, and we have assumed that this might vary between $\frac{1}{4}$ and $\frac{3}{4}$. But for every probability p within these limits (and, in fact, for every p between 0 and 1), the combined probability of the first two columns or joint events is $\frac{1}{2}$, and similarly for the last two columns. Thus there is no "least favorable distribution" distinct from the "most favorable distribution" over payoffs; for every component distribution over (Red, Black) there is the same distribution over payoffs, probability $\frac{1}{2}$ of \$100 and probability $\frac{1}{2}$ of \$0. In terms of expected utilities, $x_{\text{est}} = x_{\text{min}} = 5$, so the decision rule gives: $p.5 + (1-p)5 = 5$. This implies that you should be willing to pay \$40 for this lottery ticket, or more than you would for either of its "prizes," actions III or IV, if offered separately! The rule clearly forces you to violate Samuelson's Independence Axiom, since you prefer (III, IV; $\frac{1}{2}$) to (III, II; $\frac{1}{2}$) although you were indifferent between III and IV. Likewise, it implies that you should be indifferent between the prospect (III, IV; $\frac{1}{2}$) and the prospect (I, II; $\frac{1}{2}$), despite the fact that both I and II are preferred either to III or to IV considered separately.

here So much for the implications of the decision rule; are these results likely to be accepted as corresponding to the decision-maker's actual preferences? As we have analyzed the prospect (III, IV; $\frac{1}{2}$) above, it offers exactly the same distribution of payoffs as action I or action II; we should be disturbed if the decision rule did not assign them the same value, and surprised if a decision-maker were not, after such an analysis, indifferent between them. In comparing the lottery ticket with the two component "prizes," we note that all three alternatives offer the same "best guess" distributions over the payoffs \$100, \$0, but that the former two alternative distributions

$\#$ The utility axioms would then compel us to reconsider our earlier evaluation of the two component "prizes," actions III and IV. In comparing each of them, separately, with the lottery ticket (III, IV; $\frac{1}{2}$),

(III, IV; $\frac{1}{2}$)

Red (p)		Black ($1-p$)	
Heads ($\frac{1}{2}$)	Tails ($\frac{1}{2}$)	Heads ($\frac{1}{2}$)	Tails ($\frac{1}{2}$)
\$100	\$0	\$0	\$100

7

8
8
8
8
8
43

The probability of any one of these joint events will depend on the probability p assigned to a drawing of Red from Urn I, and we have assumed that this might vary between $1/4$ and ~~$3/4$~~ $3/4$. But for every probability p within these limits (and, in fact, for every p between 0 and 1), the probability distribution over the payoffs \$100, \$0 is precisely $1/2-1/2$. But the probability distribution over payoffs "given" that a Red ball is drawn is unambiguously $1/2-1/2$; and the probability distribution over payoffs given that a Black ball is drawn is the same, $1/2-1/2$. Thus there is no "least favorable distribution" distinct from the "most favorable distribution" over payoffs; whether Red or Black is drawn, and hence for every component distribution over (Red, Black) there is the same distribution over payoffs, probability $1/2$ of \$100 and probability $1/2$ of \$0.

1. A somewhat less "logical" grouping of the same joint outcomes of the coin toss and the drawing may make this result more obvious:

(III, IV; $1/2$)	$\frac{1}{2}p$ Heads/Red \$100	$\frac{1}{2}(1-p)$ Tails/Black \$100	$\frac{1}{2}(1-p)$ Heads/Black \$0	$\frac{1}{2}p$ Tails/Red \$0
		$\frac{1}{2}$		$\frac{1}{2}$

~~For example~~ The probability p of any one of these joint events will depend on the probability p assigned to a drawing of Red from Urn I, which we have assumed might vary between $1/4$ and $3/4$. But for every probability p within these limits (and in fact, for p every p between 0 and 1), the combined probability of the first two columns or joint events is $\frac{1}{2}$, and similarly for the last two columns. Hence there is ~~xxxxx~~ is a single, unambiguous distribution over the payoffs, \$100, \$0.

75
43
28
13
8

But even without this calculation, we could have derived ^{the} this result immediately from the observations that lie behind it if we assume that you obey Savage's Postulate 1, the complete ordering of actions. The fact that you are willing to pay the same amount for ~~an urn with~~ "100 on Red_{II}" as for "100 on Heads" with your fair coin, and the same amount for "100 on Black_{II}" (or "100 against Red_{II}") as for "100 on Tails" means that you assign the same probability to Red_{II} as to Heads; and we have already found the latter ^{probability} to be $\frac{1}{2}$. The utility calculations merely reflect these observed correspondences.

Now we consider bets with respect to Urn I, for which you have not been told the proportion of red to black balls. ~~We~~ Let us assume that you are one of those who would pay something less ~~than~~ for "100 on Red_I" or for "100 on Black_I", if ~~one of those~~ ^{contingencies} were offered you, than you would for "100 on Red_{II}". For a specific example, suppose that you follow the decision rule $[p \cdot y^0 + (1-p) y_{\min}](\chi)$ with ~~the~~ ^{1st p. 12}, and that having

observed a certain ^{1st p. 12} sample of balls ~~of~~ from ~~the~~ Urn I your "best guess" ~~discrete~~ distribution with $p = \frac{1}{2}$, y^0 is $(\frac{1}{2}, \frac{1}{2})$ and the set of reasonable distributions Y^0 ranges from $(\frac{1}{4}, \frac{3}{4})$ to $(\frac{3}{4}, \frac{1}{4})$. This would indicate a value to you (in utility) of 3.75 for the bet on Red_I, or something less than \$40 (since \$40 has utility 5). We could check to see if this rule did correspond ~~to your beliefs and preferences~~ by removing red balls and adding black in Urn II until we found you indifferent between betting on Red_I and Red_{II}; with the parameters given, this should occur with roughly 37 red and 63 black in Urn II. However, we are interested ~~here~~ only in the implications of following such a rule, not in its realism.

Finally, we offer you a lottery ticket which has the following two acts as prizes:

	Red _I	Black _I
III	\$100	\$0
IV	\$0	\$100

If your fair coin comes up Heads, you will "receive" action III, "a chance on Red_I"; if it comes up Tails, you will ~~receive~~ ^{choose} action IV may choose action

offered, alternatively, " $\frac{1}{2}$ chance on Black." Substituting the known utility numbers for the money outcomes, we assert: $p.10 + (1-p).0 = 5$, when p is your subjective probability of drawing a red ball from Urn II, so that we can calculate directly:

$$p = \frac{1}{2}.$$

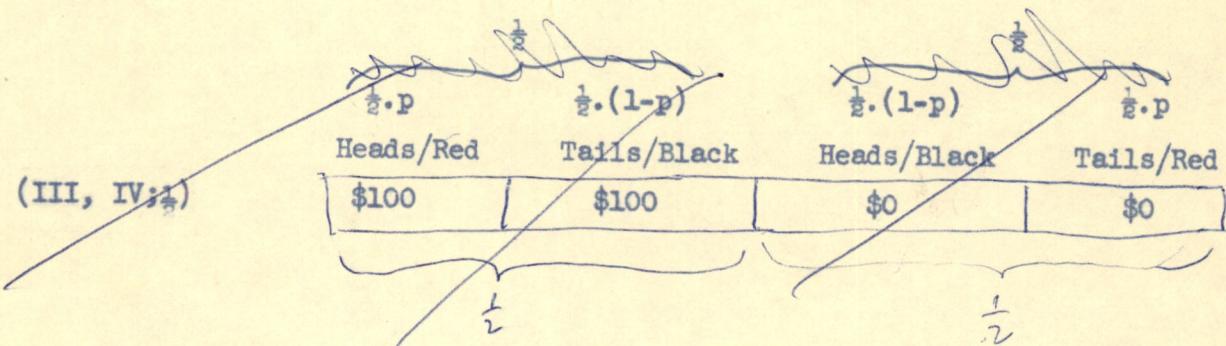
front B

Finally, we offer you the following option: if a fair coin comes up Heads, you will be offered action III below; if it comes up Tails, you will be offered action IV.

		100
	Red I	Black I
III	\$100	\$ 0
IV	\$ 0	\$100

This is a lottery ticket with actions III and IV as prizes: the prospect (III, IV; $\frac{1}{2}$). Since you are indifferent between III and IV and would pay the same amount X for either of them (X \$40, $U(X) = 3.75$) then if you were to obey the utility axioms when bets of this nature were offered as prizes in lotteries you would be indifferent between this lottery ticket and either III or IV considered separately; you should be willing to pay for this prospect the same amount X you would give, say, for III.

But consistent application of your decision rule would decree otherwise. Your payoffs now depend on the joint results of flipping a coin and drawing a ball; there are four possibilities shown below with the corresponding outcomes:



We shall further suppose that we have used this coin to calibrate your utility scale, by constructing lottery tickets with amounts of cash contingent upon the occurrence of given sequences and observing your choices among these prospects. Thus, assuming we have started by assigning arbitrarily to the money outcomes A and C the utility numbers $\frac{10}{10}$ and $\frac{0}{0}$, if we represent the outcome B by the utility number 5 this means that you have actually been observed to be indifferent between B and the prospect $(A, C; 1/2)$ (where A is contingent upon, say, the occurrence of Heads on a single flip of the fair coin); or else it means that you can reliably be expected to be indifferent, on the basis of other observed choices and your general consistency with the utility axioms in choosing among prospects. To specify a particular set of von Neumann-Morgenstern utilities, as we shall do below, is to postulate that you are a person with particular, well-defined preferences among given prospects.

Let us suppose that the outcomes A and C above are \$100 and ~~\$20~~^{\$0}, respectively, and that you happen to be indifferent between the sure possession of \$40 and a prospect offering \$100 if a fair coin comes up Heads and \$0 if it comes up Tails. We assign \$40 the utility number 5. Suppose, further, that we have found you indifferent between the following two actions, where R_{II} signifies drawing a red ball, B_{II} a black ball, from our original Urn II, known to contain exactly 50 red and 50 black balls.

	50	50	50
	R_{II}		B_{II}
I	\$100		\$0
II	\$ 40		\$40

In other words, you are willing to pay up to \$40 for "~~a~~^{\$100} chance on Red."

Similarly, we suppose that you are willing to pay the same amount if you are

event:

	E	\bar{E}
I	A	C
II	B	C

The Strong Independence Assumption then becomes the familiar proposition that action I is preferred to II, or II is preferred to I, or I and II are indifferent, as A is preferred to B, B preferred to A, or A and B are indifferent. The only difference, in this instance, from the Sure-Thing Principle is that the probability of E is assumed "known," and A, B, C are lottery tickets with known probabilities (including such "lotteries" as [A;1], [B;1], etc., otherwise known as "sure prospects").¹

In early critiques of the axioms, the counter-examples that turned up were rather an odd lot. They included the player of Russian Roulette, or the mountaineer who sought out (without publicity) the most dangerous peaks;² another is a man who would not bet a dime to a dollar that the sun will rise tomorrow.²³ It is possible to explain such behavior in other terms, but even if such persons are interpreted as violating the axioms, they do not contradict strongly the view expressed by Robert Strotz: "In my own opinion, it would be a strange man indeed who would persist in violating these precepts once he understood clearly in what way he was violating them."³⁴ These people are pretty strange. Russian

¹² Both of these could be held to violate Axiom 3:B:b; of two prospects offering the same two outcomes, they prefer one with a higher probability of the less desirable outcome.

³² He violates 3:B:c, the Continuity axiom: if A is preferred to B and B is preferred to C, there must be some prospect (A,C:p) that is preferred to B, with p less than 1.

³⁴ Robert H. Strotz, "Cardinal Utility," American Economic Review, May 1953, pp. "state of the world."

1. The close relation, in spirit, ~~but not~~ to the notion of "admissibility" or "non-domination" is evident here; but note that A, and B and C are not, in general, "consequences," but are sets of possible consequences ~~most~~ corresponding to the distinct "states of the world" that comprise the events E and \bar{E} . Neither I nor II will, in general, dominate the other; i.e. have a preferred or indifferent consequence for every

roulette players may not be crazy, but it might well be argued that a working theory of rational choice could afford to ignore them.

More recently, Maurice Allais and others have produced some examples, one of which we shall consider later, which tempt less eccentric subjects to violate the axioms. Of course, even the most pronounced adherents of the axioms expect that people will "frequently and flagrantly behave in ¹disaccord with the utility theory," just as they will violate any normative theory, or principles of logic or arithmetic. What is claimed is that it is unreasonable for them to do so, and that it should be possible through analysis and demonstration to persuade an intelligent person to see this and to "correct" his decisions, at least if he were given "a couple of years to brood over his answers."²

Let us now consider prospects that offer as "outcomes" actions whose consequences are ambiguous; in other terms, lottery tickets which offer (with "known" probabilities) other bets as prizes. We could say, "with other lottery tickets as prizes," except that I wish to reserve the terms prospect and lottery ticket for bets the probabilities of whose outcomes are stipulated to be known definitely to the subject: i.e., the outcomes of which are contingent upon events with respect to which he obeys the Savage axioms.

We shall assume that such events exist. Specifically, let us assume that we have discovered a coin you are convinced is fair; you act "as if" you assigned definite probabilities to every finite sequence of heads and

tails, and the same probability to every sequence of the same length. ³ *This dispelling my earlier, pessimistic conjectures about the ease of finding such a coin, or, I should say, such a subject; again we see the advantages of hypothetical experiments as a basis for discussion.*

¹ Savage, op. cit., p. 100.

² Samuelson, op. cit.,